NVM Express Technical Errata

Errata ID	**004**
Change Date	**5/7/2013**
Affected Spec Ver.	**NVM Express 1.0 and 1.1**
Corrected Spec Ver.	

Submission info

Name	Company	Date
Amber Huffman	Intel	2/7/2013
John Carroll	Intel	2/7/2013
Matthew Wilcox	Intel	2/7/2013
Peter Onufryk	IDT	2/7/2013
Kevin Marks	Dell	2/21/2013

Section 3 is unclear on controller register access requirements. Clarifications are made on register access rules, noting access size, alignment requirements, and avoiding multiple registers being accessed at the same time.

Description of the specification technical flaw:

Modify the first two paragraphs of section 3 as shown below:

Controller registers are located in the MLBAR/MUBAR registers (PCI BAR0 and BAR1) that shall be mapped to a memory space that supports in-order access and variable access widths. For many computer architectures, specifying the memory space as uncacheable produces this behavior. The host shall not issue locked accesses. The host shall access registers in their native width or aligned 32-bit accesses. Violation of either of these host requirements results in undefined behavior. ~~Locked accesses are not supported and registers shall be accessed in their native width (16-bit registers should be accessed as a 16-bit value, 32-bit registers should be accessed as a 32-bit value, and 64-bit registers should be accessed as a 64-bit value). For example, the 32-bit Submission Queue 0 Tail Doorbell register should be accessed as a 32-bit value. If software cannot access the register in its native width due to limitations of the platform, then it should access the register in order starting with lower bytes first. For example, accessing a 64-bit register using 16-bit accesses would result in an access for bits 15:00, following by an access for bits 31:16, followed by an access for bits 47:32, and concluded with an access for bits 63:48.~~

Accesses that target any portion of two or more registers are not supported.

~~Unaligned register accesses are not supported and are handled at the PCI Express layer with a PCI Express completer abort status, indicating a violation of the programming model.~~

Disposition log

2/7/2013	Erratum captured.
2/21/213	Discussed at the technical workgroup – clarified wording.
3/21/2013	Added clarification for accessing 2+ registers.
4/2/2013	Clarified unaligned doorbell write wording.
4/18/2013	Clarified that accesses shall be native width or 32-bit aligned.
5/7/2013	Clarified the wording to indicate host requirements and if not followed, undefined behavior.
6/7/2013	Erratum ratified.